Semantic Diversity versus Visual Diversity in Visual Dictionaries

Otávio A. B. Penatti, Sandra Avila, *Member, IEEE,* Eduardo Valle, Ricardo da S. Torres, *Member, IEEE*

arXiv:1511.06704v1 [cs.CV] 20 Nov 2015

***Abstract*—Visual dictionaries are a critical component for image classification/retrieval systems based on the bag-of-visual-words (BoVW) model. Dictionaries are usually learned without supervision from a training set of images sampled from the collection of interest. However, for large, general-purpose, dynamic image collections (e.g., the Web), obtaining a representative sample in terms of semantic concepts is not straightforward. In this paper, we evaluate the impact of semantics in the dictionary quality, aiming at verifying the importance of semantic diversity in relation visual diversity for visual dictionaries. In the experiments, we vary the amount of classes used for creating the dictionary and then compute different BoVW descriptors, using multiple codebook sizes and different coding and pooling methods (standard BoVW and Fisher Vectors). Results for image classification show that as visual dictionaries are based on low-level visual appearances, visual diversity is more important than semantic diversity. Our conclusions open the opportunity to alleviate the burden in generating visual dictionaries as we need only a visually diverse set of images instead of the whole collection to create a good dictionary.**

I. Introduction

Among the most effective representations for image classification, many are based on visual dictionaries, in the so-called Bag of (Visual) Words (BoW or BoVW). The BoVW model brings to Multimedia Information Retrieval many intuitions from Textual Information Retrieval [1], mainly the idea of using occurrence statistics on the contents of the documents (words or local patches), while ignoring the document structure (phrases or geometry).

The first challenge in translating the BoVW model from textual documents to images is that the concept of a "visual word" is much less straightforward than that of a textual word. In order to create such "words," one employs a visual dictionary (or codebook), which organizes image local patches into groups according to their appearance. This description fits the overwhelming majority of works in the BoVW literature, in which the "visual words" are exclusively appearance-based, and do not incorporate any semantic information (e.g., from class labels). This is in stark contrast to textual words, which are semantically rich. Some works incorporate semantic information from the class labels into the visual dictionary [2]–[4], but the benefit brought by those schemes is, so far, small compared to their cost, especially for large-scale collections.

The BoVW model is a three-tiered model based on: (1) the extraction of low-level local features from the images, (2) the aggregation of those local features into mid-level BoVW feature vectors (guided by the visual dictionary), (3) the use of the BoVW feature vectors as an input for a classifier (often Support Vector Machines, or SVM). The most traditional way to create the visual dictionary is to sample low-level local features from a learning set of images, and to employ unsupervised learning (often k-means clustering) to find the k vectors that best represent the learning sample.

Aiming at representativeness, the dictionary is almost always learned on a training set sampled from the same collection that will be used for training and testing the classifier. In the closed datasets [5] popularly used in the literature, the amount of images is fixed, therefore no new content is added after the dictionary is created. However, in a large-scale dynamic scenario, like the Web, images are constantly inserted and deleted. Updating the visual dictionary would imply updating all the BoVW feature vectors, which is unfeasible.

It might be possible to solve this problem if we remind that the term "visual dictionary" is somewhat a misnomer, because it is not concerned with semantic information. The visual dictionaries we are interested in this work ignore the class labels and consider only the appearance-based low-level features. Provided that the selected sample represents well enough that low-level feature space (i.e., being diverse in terms of appearances), the dictionary obtained will be sufficiently accurate, even if it does not represent all semantic variability.

We have seen recently an exponential growth of deep learning, mainly convolutional (neural) networks or ConvNets, for visual recognition [6]–[8]. However, some recent works [9], [10] are joining both worlds: ConvNets and techniques based on visual dictionaries, like Fisher Vectors. Therefore, visual codebooks still tend to remain in the spotlight for some years.

This paper evaluates the hypothesis that the quality of a visual dictionary depends only on the visual diversity of the samples used. The experiments evaluate the impact of the semantic diversity in relation to the visual diversity, varying the amount of classes used as basis for the dictionary creation. Results point to the importance of visual variability regardless of semantics. We show that dictionaries created on few classes (i.e., low semantic variability) are able to produce good image representations that are as good as representations that use dictionaries based on the whole set of classes. To the best of our knowledge, there is no similar objective evaluation of visual dictionaries in the literature.

II. Related work

We start this section with a review of the existing art on dataset dependency of dictionaries. Then, we detail the

O. A. B. Penatti is with the SAMSUNG Research Institute, Campinas, Brazil (e-mail: o.penatti@samsung.com).

S. Avila and E. Valle are with the School of Electrical and Computing Engineering, University of Campinas (Unicamp), Campinas, Brazil.

R. da S. Torres is with the Institute of Computing, University of Campinas (Unicamp), Campinas, Brazil.

representation based on visual dictionaries. This work focuses on visual dictionaries created by unsupervised learning, since there are several challenges in scaling-up supervised dictionaries. However, at the end of this section, we also briefly discuss how supervised dictionaries contrast to our experiments.

A. Dictionaries and dataset dependency

There is scarce literature on the impact of the sampling choice on the quality of visual dictionaries. According to Perronnin and Dance [11], dictionaries trained on one set of categories can be applied to another set without any significant loss in performance. They remark that the Fisher kernels are fairly insensitive to the quality of the generative models and therefore that the same visual dictionary can be used for different category sets. However, they do not evaluate the impact of semantics in the dictionary quality.

To tackle the problem of dataset dependency for the visual dictionaries, Liao et al. [12] transform BoVW vectors derived from one visual dictionary to make them compatible with another dictionary. By solving least squares, the authors showed that the cross-codebook method is comparable to the within-codebook one, when the two codebooks have the same size.

Zhou et al. [13] highlight many of the problems that we discuss here: the problem of dataset dependency for the visual dictionaries, the high computational cost for creating dictionaries, and the problem of update inefficiency. Although they do not provide experiments for confirming that, they propose a scalar quantization which is independent of datasets.

B. Image representations based on visual dictionaries

The bag-of-visual-words (BoVW) model [1], [14] depends on the visual dictionary (also called visual codebook or visual vocabulary), which is used as basis for encoding the low-level features. BoVW descriptors aim to preserve the discriminating power of local descriptions while pooling those features into a single feature vector [4].

The pipeline for computing a BoVW descriptor is divided into (i) dictionary generation and (ii) BoVW computation. The first step can be performed off-line, that is, it is based on the training images only. And the second step computes the BoVW using the dictionary created in the first step.

To generate the dictionary, one usually samples or clusters the low-level features of the images. Those can be based on interest-point detectors or on dense sampling in a regular grid [15]–[17]. From each of the points sampled from an image, a local feature vector is computed, with SIFT [18] being a usual choice. Those feature vectors are then either clustered (often based on k-means) or randomly sampled in order to obtain a number of vectors to compose the visual dictionary [16], [19], [20]. This can be understood as a feature space quantization.

After creating the dictionary, the BoVW vectors can be computed. For that, the original low-level feature vectors need to be coded according to the dictionary – *coding* step. This can be simply performed by assigning one or more of the visual words to each of the points in an image. Popular coding approaches are based on *hard* or *soft* assignment [21], [22], although, several possibilities exist [23]–[27].

The *pooling* step takes place after the coding step. One can simply average the codes obtained over the entire image (average pooling), can consider the maximum visual word activation (max pooling [4], [24]), or can use several other existing pooling strategies [21], [28]–[30]

Recently, Fisher Vectors [11], [31] have gained attention due to their high accuracy for image classification [17], [32], [33]. Fisher Vectors extend BoVW by encoding the average first- and second-order differences between the descriptors and visual words. In [31], Perronnin et al. proposed modifications over the original framework [11] and showed that they could boost the accuracy of image classifiers. A recent comparison of coding and pooling strategies is presented in [34].

C. Supervised visual dictionaries

Many supervised approaches have been proposed to construct discriminative visual dictionaries that explicitly incorporate category-specific information [4], [35]–[39]. In [40], the visual dictionary is trained in unsupervised and supervised learning phases. The visual dictionary may even be generated by manually labeling image patches with a semantic label [41]–[43]. Other more sophisticated techniques have been adapted to learn the visual dictionary, such as restricted Boltzmann machines.

The main problem of those supervised dictionaries is that we need to know beforehand – when the dictionary is being created – all the possible classes in the dataset. For dynamic environments, this is extremely challenging. Another issue is the difficulty of finding enough labeled data that be representative of very large collections. Finally, many supervised approaches for dictionary learning rely on very costly optimizations, for a small increase in accuracy. For those reasons, we focus on dictionaries created by unsupervised learning, which also represent the majority of works in the literature.

III. Experimental setup

We evaluated the impact of semantics in the dictionary quality, by varying the number of classes used as basis for the dictionary creation. We tested different codebook sizes and different coding and pooling strategies. It is important to evaluate those parameters as there could exist effects that appear only on specific configurations.

We report the results using the PASCAL VOC 2007 dataset [44], a challenging image classification benchmark that contains significant variability in terms of object size, orientation, pose, illumination, position and occlusion [5]. This dataset consists of 20 visual object classes in realistic scenes, with a total of 9,963 images categorized in macro-classes and sub-classes as person (*person*), animal (*bird*, *cat*, *cow*, *dog*, *horse*, *sheep*), vehicle (*aeroplane*, *bicycle*, *boat*, *bus*, *car*, *motorbike*, *train*), and indoor objects (*bottle*, *chair*, *dinning table*, *potted plant*, *sofa*, *tv/monitor*). Those images are split into three subsets: training (2,501 images), validation (2,510 images) and test (4,952 images). Our experimental results are obtained on 'train+val'/test sets.

Before extracting features, all images were resized to have at most 100 thousand of pixels (if larger), as proposed in [45]. As low-level features, we have extracted SIFT descriptors [46] from 24 $\times$ 24 patches on dense regular grid every 4 pixels at 5 scales. The dimensionality of SIFT was reduced from 128 to 64 by using principal component analysis (PCA). We have used one million descriptors randomly selected and uniformly divided among the training images (i.e., the same number of descriptors per image) as basis for PCA and also for dictionary generation. As the random selection could also impact the results, two sets of one million descriptors were used.

We used two different coding/pooling methods and their visual dictionaries were created accordingly. For BoVW [1] (obtained with hard assignment and average pooling), we apply k-means clustering with Euclidean distance over the sets of one million descriptors explained above. We have varied the number of visual words in 1,024, 2,048, 4,096, and 8,192.

For Fisher Vectors [31], the descriptor distribution is modeled using a Gaussian mixture model, whose parameters are also trained over the sets of one million descriptors just mentioned, using an expectation maximization algorithm. The dictionary sizes were 128, 256, 512, 1,024, and 2,048 words.

We used a classification protocol (one-versus-all) with support vector machines (SVM) with the most appropriate configuration for each coding/pooling method: RBF kernel for BoVW and linear kernel for Fisher Vectors [47]. Grid search was performed to optimize SVM parameters. The classification performance was measured by the average precision (AP).

One of the most important parameters in the evaluation was the variation in the number of image classes used for dictionary generation. We used 18 different samples from the VOC 2007 dataset. The initial selection of samples was based on selecting individual classes, varying from 1, 2, 5, 10 to all the 20 classes. However, as VOC 2007 is multilabel (i.e., an image can belong to more than one class), when selecting points from images of class *cat*, for instance, we probably also select points of class *sofa*. Therefore, we show our results in terms of the real number of classes used to create each sample. We count the number of classes that appear in each sample, even if only 1 image of a class is used. We also computed a measure that considers the frequency of occurrence of classes, which we called *semantic diversity*, better estimating the real amount of semantics in the samples. The semantic diversity S is measured by the following equation:

$$S = \sum_{i=1}^{N} \frac{C_i + {C'}_i}{C_i}, \qquad (1)$$

where N is the number of classes selected, C_i is the number of images in the class i and ${C'}_i$ is the number of images in the class i that also appear in other classes. A summary of the samples used is shown in Table I.

It is important to highlight that the selection of classes (dictionary classes) is used only for creating the dictionary. The results (target classes) are reported always considering the whole VOC 2007 dataset.

For analyzing the results, we took the logit of the average precisions in order to have measures that behave more linearly (i.e., that are more additive and homogeneous throughout the scale). Thus each average precision was transformed using the function: $logit(x) = -10 \times \log_{10}(\frac{x}{1-x})$. Then, in order to get the global effect of the increasing number of dictionary classes, we have aggregated the average precisions for all target classes. To make the average precision of different target classes commensurable, we have "erased" the effect of the target classes themselves by subtracting the average and dividing the standard deviation obtained for all logit values for that target class. This procedure, commonly employed in statistical meta-analyses, is necessary to obtain an aggregated effect because some target classes are much harder to classify than others. Let us call those values the standardized logits of the average precisions, which will be used in our analyzes.

id	S	N_c	N_i	initial classes selected
1	1.49	18	421	dog
2	1.56	18	713	car
3	2.24	15	445	chair
4	3.14	11	327	cow, bus
5	3.24	12	561	cat, sofa
6	3.57	13	530	dining, bird
7	3.93	15	501	tv, motorbike
8	4.11	17	731	horse, chair
9	7.16	19	1560	cat, bird, train, plant, dog
10	8.75	20	1638	car, sofa, bottle, bicycle, train
11	8.82	20	1306	motorbike, sheep, horse, aeroplane, chair
12	8.84	20	2477	cow, boat, table, bus, person
13	9.78	18	2602	bus, chair, bicycle, sofa, person
14	16.28	20	3478	cat, horse, person, sheep, train, tv/monitor, bird, cow, sofa, motorbike
15	16.65	20	3358	aeroplane, bird, plant, chair, horse, car, sofa, dog, boat, bicycle
16	16.87	20	2491	horse, cow, bus, aeroplane, table, chair, boat, bicycle, dog, cat
17	16.88	20	3749	tv/monitor, car, motorbike, bottle, cow, bird, train, plant, sheep, person
18	18.02	20	2366	bus, cow, boat, table, aeroplane, bicycle, bird, train, dog, bottle

TABLE I

SUMMARY OF THE SAMPLES USED AS BASIS FOR THE DICTIONARY CREATION IN ALL OF THE EXPERIMENTS IN THIS SECTION. S REPRESENTS THE SEMANTIC DIVERSITY COMPUTED BY EQUATION 1, N_c IS THE NUMBER OF CLASSES THAT APPEAR AT LEAST ONCE IN EACH SAMPLE AND N_i IS THE NUMBER OF IMAGES IN THE SAMPLE.

IV. RESULTS

Each experiment configuration explained in Section III generated a set of 40 points (20 target classes and 2 runs, each one with a 1 million feature set). After plotting the standardized logits of the average precisions, we have computed the linear regression aiming at determining the variation in results when varying from low to high semantics (increasing number of dictionary classes).

Figures 1 and 2 show the results for BoVW. In Figure 1, all the dictionary configurations are shown together while in Figure 2, each dictionary is shown separately. We can see that, as semantics increase in the dictionary (i.e., more classes are considered during dictionary creation), there is almost no increase in average precision. This is numerically visible by the very small values of the red line slope α. We can also attest with high confidence (small p-value) that semantics has low importance for the dictionary quality.

For Fisher Vectors, the results are presented in Figures 3 and 4. Similarly to BoVW, semantics has few impact in the accuracy of Fisher Vectors. Although in some specific configurations, i.e., dictionaries of 1024 and 2048 (Figures 4-*d* and *j*) the regression curve has a larger slope, its coefficient is still very small.

Therefore, we can conclude that the impact of semantics in the dictionary quality is very low. Although the samples used for creating the dictionary can still be very poor in terms of

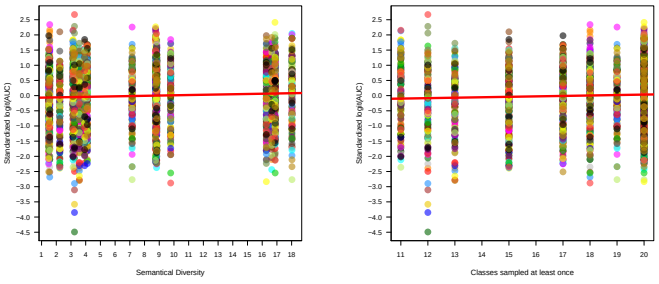

(a) α = 0.0085 (p-value=0.0071) (b) α = 0.0146 (p-value=0.0197)

Fig. 1. Standardized average precision of each class versus semantic variability considering all the runs of every BoVW configuration (all codebook sizes: 1024, 2048, 4096, 8192). The red line corresponds to the linear regression. The coefficient α (slope) is very small, showing that semantics has a very low impact on the codebook quality. Each of the 20 classes has a different dot color. In (a) the x-axis shows the semantic diversity, while in (b), the x-axis corresponds to the number of classes in each sample used to create the dictionary. Semantic diversity and number of classes are shown in Table I.

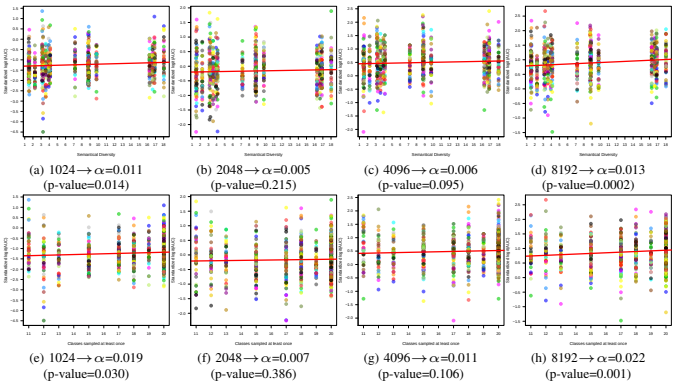

(a) 1024→α=0.011 (p-value=0.014) (b) 2048→α=0.005 (p-value=0.215) (c) 4096→α=0.006 (p-value=0.095) (d) 8192→α=0.013 (p-value=0.0002)

(e) 1024→α=0.019 (p-value=0.030) (f) 2048→α=0.007 (p-value=0.386) (g) 4096→α=0.011 (p-value=0.106) (h) 8192→α=0.022 (p-value=0.001)

Fig. 2. Similar graphs of Figure 1 showing that even when we analyze the BoVW results per dictionary, we can see a very low impact of semantics in the precision. The first row has the plots versus semantic diversity and, the second row, versus the number of classes in each sample.

semantics (few classes), they can be rich in terms of visual diversity, which is enough for creating a good dictionary. As the low-level descriptor (SIFT, in our experiments) captures image local properties and not semantics, the fast dictionary generalization occurs if we have a set of images rich enough in terms of textures, which will cover a good portion of the feature space without requiring the use of all image classes. We can also conclude that the dictionary size and the coding/pooling have a minor impact in those conclusions, as the observations above were possible regardless of the dictionary size and regardless of the coding/pooling method.

V. Discussion

In this paper, we evaluate the impact of semantic diversity in the quality of visual dictionaries. The experiments conducted show that dictionaries based on a subset of a collection may still provide good performance, provided that the selected sample is visually diverse. We showed experimentally that the impact of semantics in the dictionary quality is very small: classification results of dictionaries based on a sample with low or high semantic variability were similar. Therefore, we can point out that visual variability is more important than

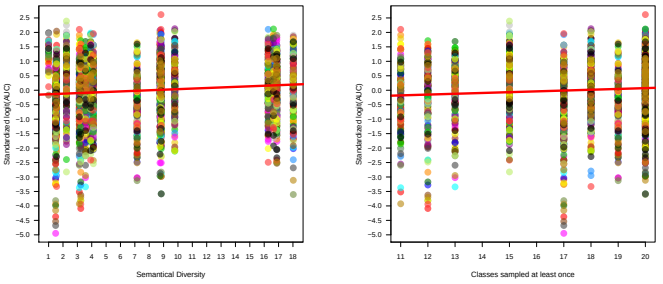

(a) α = 0.0196 (p-value=8.3e-12) (b) α = 0.0279 (p-value=8.4e-7)

Fig. 3. For Fisher Vectors, similarly to BoVW (see Figure 1), we can affirm with high confidence (very small p-value) that semantics has a very low impact on the codebook quality (coefficient α is very small).

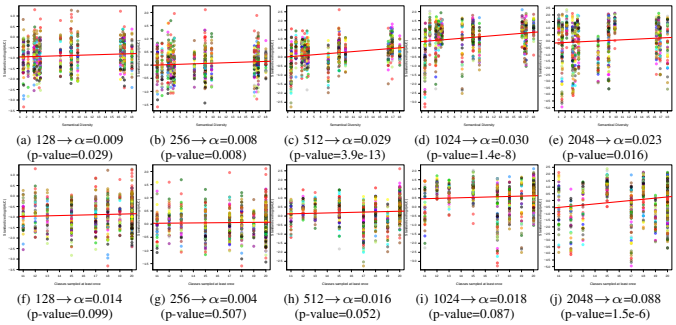

(a) 128→α=0.009 (p-value=0.029) (b) 256→α=0.008 (p-value=0.008) (c) 512→α=0.029 (p-value=3.9e-13) (d) 1024→α=0.030 (p-value=1.4e-8) (e) 2048→α=0.023 (p-value=0.016)

(f) 128→α=0.014 (p-value=0.099) (g) 256→α=0.004 (p-value=0.507) (h) 512→α=0.016 (p-value=0.052) (i) 1024→α=0.018 (p-value=0.087) (j) 2048→α=0.088 (p-value=1.5e-6)

Fig. 4. Analogously to Figure 2, Fisher Vectors have low variation in accuracy when varying the amount of semantics used for creating the dictionaries.

semantic variability when selecting the sample to be used as source of features for the dictionary creation.

The scenarios that can benefit from our conclusions are the ones in which the dataset is dynamic (images are inserted and removed during the application use) and/or in which one do not have access to the whole dataset when the dictionary is created. Those scenarios are important because they reflect any application dealing with web-like collections.

As a general recommendation for creating visual dictionaries in dynamic environments, we say that one must have a diverse set of images in terms of visual appearances and not necessarily in terms of classes. Those findings open the opportunity to greatly alleviate the burden in generating the dictionary, since, at least for general-purpose datasets, we show that the dictionaries do not have to take into account the entire collection, and may even be based on a small collection of visually diverse images.

In future work, we would like to explore if those findings also work on special-purpose datasets, like in quality control of industry images or in medical images, in which we suspect that a specific dictionary might have a stronger effect. We also would like to investigate possibilities for visualizing the codebooks, aiming at better understanding how different dictionaries cover the feature space.

Acknowledgment

Thanks to Fapesp (grant number 2009/10554-8), CAPES, CNPq, Microsoft Research, AMD, and Samsung.

REFERENCES

[1] J. Sivic and A. Zisserman, "Video google: a text retrieval approach to object matching in videos," in *International Conference on Computer Vision*, vol. 2, 2003, pp. 1470–1477.

[2] F. Perronnin, C. Dance, G. Csurka, and M. Bressan, "Adapted vocabularies for generic visual categorization," in *European Conference on Computer Vision*, 2006, pp. 464–475.

[3] S. Lazebnik and M. Raginsky, "Supervised learning of quantizer codebooks by information loss minimization," *Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 7, pp. 1294–1309, 2009.

[4] Y.-L. Boureau, F. Bach, Y. LeCun, and J. Ponce, "Learning mid-level features for recognition," *Conference on Computer Vision and Pattern Recognition*, pp. 2559–2566, 2010.

[5] A. Torralba and A. A. Efros, "Unbiased look at dataset bias," in *Conference on Computer Vision and Pattern Recognition*, 2011, pp. 1521–1528.

[6] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Neural Information Processing Systems*, 2012, pp. 1106–1114.

[7] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun, "Overfeat: Integrated recognition, localization and detection using convolutional networks," in *International Conference on Learning Representations*, 2014.

[8] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, "Cnn features off-the-shelf: An astounding baseline for recognition," in *Conference on Computer Vision and Pattern Recognition Workshop*, 2014, pp. 512–519.

[9] F. Perronnin and D. Larlus, "Fisher vectors meet neural networks: A hybrid classification architecture," in *Conference on Computer Vision and Pattern Recognition*, June 2015, pp. 3743–3752.

[10] B. Klein, G. Lev, G. Sadeh, and L. Wolf, "Associating neural word embeddings with deep image representations using fisher vectors," in *Conference on Computer Vision and Pattern Recognition*, June 2015, pp. 4437–4446.

[11] F. Perronnin and C. Dance, "Fisher kernels on visual vocabularies for image categorization," in *Conference on Computer Vision and Pattern Recognition*, 2007, pp. 1–8.

[12] S. Liao, X. Li, and X. Du, "Cross-codebook image classification," in *Pacific-Rim Conference on Multimedia*, 2013, pp. 497–504.

[13] W. Zhou, Y. Lu, H. Li, and Q. Tian, "Scalar quantization for large scale image search," in *ACM Multimedia*, 2012, pp. 169–178.

[14] G. Csurka, C. Bray, C. Dance, and L. Fan, "Visual categorization with bags of keypoints," in *Workshop on Statistical Learning in Computer Vision, ECCV*, 2004, pp. 1–22.

[15] K. E. A. van de Sande, T. Gevers, and C. G. M. Snoek, "Evaluating color descriptors for object and scene recognition," *Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 9, pp. 1582–1596, 2010.

[16] F. Jurie and B. Triggs, "Creating efficient codebooks for visual recognition," in *International Conference on Computer Vision*, vol. 1, Oct. 2005, pp. 604–610.

[17] K. Chatfield, V. Lempitsky, A. Vedaldi, and A. Zisserman, "The devil is in the details: an evaluation of recent feature encoding methods," in *British Machine Vision Conference*, 2011, pp. 76.1–76.12.

[18] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.

[19] V. Viitaniemi and J. Laaksonen, "Experiments on selection of codebooks for local image feature histograms," in *International Conference on Visual Information Systems: Web-Based Visual Information Search and Management*, 2008, pp. 126–137.

[20] J. P. Papa and A. Rocha, "Image categorization through optimum path forest and visual words," in *International Conference on Image Processing*, 2011, pp. 3525–3528.

[21] L. Liu, L. Wang, and X. Liu, "In defense of soft-assignment coding," in *International Conference on Computer Vision*, 2011, pp. 1–8.

[22] J. C. van Gemert, C. J. Veenman, A. W. M. Smeulders, and J.-M. Geusebroek, "Visual word ambiguity," *Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, pp. 1271–1283, 2010.

[23] K. Yu, T. Zhang, and Y. Gong, "Nonlinear learning using local coordinate coding," in *Neural Information Processing Systems*, 2009, pp. 2223–2231.

[24] J. Yang, K. Yu, Y. Gong, and T. Huang, "Linear spatial pyramid matching using sparse coding for image classification," in *Conference on Computer Vision and Pattern Recognition*, 2009, pp. 1794–1801.

[25] J. Wang, J. Yang, K. Yu, F. Lv, T. Huang, and Y. Gong, "Locality-constrained linear coding for image classification," in *Conference on Computer Vision and Pattern Recognition*, 2010, pp. 3360–3367.

[26] G. L. Oliveira, E. R. Nascimento, A. W. Vieira, and M. F. M. Campos, "Sparse spatial coding: A novel approach for efficient and accurate object recognition," in *International Conference on Robotics and Automation*, 2012, pp. 2592–2598.

[27] S. Avila, N. Thome, M. Cord, E. Valle, and A. de A. Araújo, "Pooling in image representation: the visual codeword point of view," *Computer Vision and Image Understanding*, vol. 117, no. 5, pp. 453–465, 2013.

[28] S. Lazebnik, C. Schmid, and J. Ponce, "Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories," in *Conference on Computer Vision and Pattern Recognition*, vol. 2, 2006, pp. 2169–2178.

[29] J. Feng, B. Ni, Q. Tian, and S. Yan, "Geometric lp-norm feature pooling for image classification," in *Conference on Computer Vision and Pattern Recognition*, 2011, pp. 2609–2704.

[30] O. A. B. Penatti, F. B. Silva, E. Valle, V. Gouet-Brunet, and R. d. S. Torres, "Visual word spatial arrangement for image retrieval and classification," *Pattern Recognition*, vol. 47, no. 2, pp. 705–720, 2014.

[31] F. Perronnin, J. Sánchez, and T. Mensink, "Improving the Fisher Kernel for Large-Scale Image Classification," in *European Conference on Computer Vision*, 2010, pp. 143–156.

[32] T. M. Jorge Sánchez, Florent Perronnin and J. Verbee, "Image classification with the fisher vector: Theory and practice," *International Journal of Computer Vision*, vol. 105, no. 3, pp. 222–245, 2013.

[33] Y. Huang, Z. Wu, L. Wang, and T. Tan, "Feature coding in image classification: A comprehensive study," *Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 3, pp. 493–506, 2014.

[34] P. Koniusz, F. Yan, and K. Mikolajczyk, "Comparison of mid-level feature coding approaches and pooling strategies in visual concept detection," *Computer Vision and Image Understanding*, vol. 117, no. 5, pp. 479–492, 2013.

[35] J. Winn, A. Criminisi, and T. Minka, "Object categorization by learned universal visual dictionary," in *International Conference on Computer Vision*, vol. 2, Oct. 2005, pp. 1800–1807.

[36] F. Moosmann, B. Triggs, and F. Jurie, "Fast Discriminative Visual Codebooks using Randomized Clustering Forests," in *Neural Information Processing Systems*, 2007, pp. 985–992.

[37] F. Perronnin, "Universal and adapted vocabularies for generic visual categorization," *Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 7, pp. 1243–1256, 2008.

[38] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Supervised dictionary learning," in *Neural Information Processing Systems*, vol. 21, 2009, pp. 1033–1040.

[39] S. Lazebnik and M. Raginsky, "Supervised learning of quantizer codebooks by information loss minimization," *Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 7, pp. 1294–1309, 2009.

[40] H. Goh, N. Thome, M. Cord, and J.-H. Lim, "Unsupervised and supervised visual codes with restricted boltzmann machines," in *European Conference on Computer Vision*, 2012, pp. 298–311.

[41] J. C. van Gemert, J.-M. Geusebroek, C. J. Veenman, C. G. M. Snoek, and A. W. M. Smeulders, "Robust Scene Categorization by Learning Image Statistics in Context," in *Conference on Computer Vision and Pattern Recognition Workshop*, 2006.

[42] J. Vogel and B. Schiele, "Semantic modeling of natural scenes for content-based image retrieval," *International Journal of Computer Vision*, vol. 72, no. 2, pp. 133–157, 2007.

[43] J. Liu, Y. Yang, and M. Shah, "Learning semantic visual vocabularies using diffusion distance," in *Conference on Computer Vision and Pattern Recognition*, 2009, pp. 461–468.

[44] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results," http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html.

[45] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid, "Good practice in large-scale learning for image classification," *Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 3, pp. 507–520, 2014.

[46] A. Vedaldi and B. Fulkerson, "VLFeat - An open and portable library of computer vision algorithms," in *ACM Multimedia*, 2010, pp. 1469–1472.

[47] D. Picard, N. Thome, and M. Cord, "JKernelMachines: A simple framework for kernel machines," *Journal of Machine Learning Research*, vol. 14, pp. 1417–1421, 2013.